

Network Working Group
Request for Comments: 5330
Category: Standards Track

JP. Vasseur, Ed.
Cisco Systems, Inc
M. Meyer
BT
K. Kumaki
KDDI R&D Labs
A. Bonda
Telecom Italia
October 2008

A Link-Type sub-TLV to Convey the Number of
Traffic Engineering Label Switched Paths Signalled with
Zero Reserved Bandwidth across a Link

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

Several Link-type sub-Type-Length-Values (sub-TLVs) have been defined for Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) in the context of Multiprotocol Label Switching (MPLS) Traffic Engineering (TE), in order to advertise some link characteristics such as the available bandwidth, traffic engineering metric, administrative group, and so on. By making statistical assumptions about the aggregated traffic carried onto a set of TE Label Switched Paths (LSPs) signalled with zero bandwidth (referred to as "unconstrained TE LSP" in this document), algorithms can be designed to load balance (existing or newly configured) unconstrained TE LSP across a set of equal cost paths. This requires knowledge of the number of unconstrained TE LSPs signalled across a link. This document specifies a new Link-type Traffic Engineering sub-TLV used to advertise the number of unconstrained TE LSPs signalled across a link.

Table of Contents

| | |
|-----------------------------------|---|
| 1. Introduction | 2 |
| 2. Terminology | 3 |
| 2.1. Requirements Language | 4 |
| 3. Protocol Extensions | 4 |
| 3.1. IS-IS | 4 |
| 3.2. OSPF | 4 |
| 4. Elements of Procedure | 5 |
| 5. IANA Considerations | 5 |
| 6. Security Considerations | 5 |
| 7. Acknowledgements | 6 |
| 8. References | 6 |
| 8.1. Normative References | 6 |
| 8.2. Informative References | 6 |

1. Introduction

It is not uncommon to deploy MPLS Traffic Engineering for the sake of fast recovery, relying on a local protection recovery mechanism such as MPLS TE Fast Reroute (see [RFC4090]). In this case, a deployment model consists of deploying a full mesh of TE LSPs signalled with zero bandwidth (also referred to as unconstrained TE LSP in this document) between a set of LSRs (Label Switching Routers) and protecting these TE LSPs against link, SRLG (Shared Risk Link Group), and/or node failures with pre-established backup tunnels. The traffic routed onto such unconstrained TE LSPs simply follows the IGP shortest path, but is protected with MPLS TE Fast Reroute. This is because the TE LSP computed by the path computation algorithm (e.g., CSPF) will be no different than the IGP (Interior Gateway Protocol) shortest path should the TE metric be equal to the IGP metric.

When a reoptimization process is triggered for an existing TE LSP, the decision on whether to reroute that TE LSP onto a different path is governed by the discovery of a lower cost path satisfying the constraints (other metrics, such as the percentage of reserved bandwidth or the number of hops, can also be used). Unfortunately, metrics such as the path cost or the number of hops may be ineffective in various circumstances. For example, in the case of a symmetrical network with ECMPs (Equal Cost Multi-Paths), if the network operator uses unconstrained TE LSP, this may lead to a poorly load balanced traffic; indeed, several paths between a source and a destination of a TE LSP may exist that have the same cost, and the reservable amount of bandwidth along each path cannot be used as a tie-breaker.

By making statistical assumptions about the aggregated traffic carried by a set of unconstrained TE LSPs, algorithms can be designed to load balance (existing or newly configured) unconstrained TE LSPs across a set of equal cost paths. This requires knowledge of the number of unconstrained TE LSPs signalled across each link.

Note that the specification of load balancing algorithms is outside the scope of this document and is referred to for the sake of illustration of the motivation for gathering such information.

Furthermore, the knowledge of the number of unconstrained TE LSPs signalled across each link can be used for other purposes -- for example, to evaluate the number of affected unconstrained TE LSPs in case of a link failure.

A set of Link-type sub-TLVs have been defined for OSPF and IS-IS (see [RFC3630] and [RFC5305]) in the context of MPLS Traffic Engineering in order to advertise various link characteristics such as the available bandwidth, traffic engineering metric, administrative group, and so on. As currently defined in [RFC3630] and [RFC5305], the information related to the number of unconstrained TE LSPs is not available. This document specifies a new Link-type Traffic Engineering sub-TLV used to indicate the number of unconstrained TE LSPs signalled across a link.

Unconstrained TE LSPs that are configured and provisioned through a management system MAY be omitted from the count that is reported.

2. Terminology

Terminology used in this document:

CSPF: Constrained Shortest Path First

IGP : Interior Gateway Protocol

LSA: Link State Advertisement

LSP: Link State Packet

MPLS: Multiprotocol Label Switching

LSR: Label Switching Router

SRLG: Shared Risk Link Group

TE LSP: Traffic Engineering Label Switched Path

Unconstrained TE LSP: A TE LSP signalled with a bandwidth equal to 0

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. Protocol Extensions

Two Unconstrained TE LSP Count sub-TLVs are defined that specify the number of TE LSPs signalled with zero bandwidth across a link.

3.1. IS-IS

The IS-IS Unconstrained TE LSP Count sub-TLV is OPTIONAL and MUST NOT appear more than once within the extended IS reachability TLV (type 22) specified in [RFC5305] or the Multi-Topology (MT) Intermediate Systems TLV (type 222) specified in [RFC5120]. If a second instance of the Unconstrained TE LSP Count sub-TLV is present, the receiving system MUST only process the first instance of the sub-TLV.

The IS-IS Unconstrained TE LSP Count sub-TLV format is defined below:

Type (1 octet): 23

Length (1 octet): 2

Value (2 octets): number of unconstrained TE LSPs signalled across the link.

3.2. OSPF

The OSPF Unconstrained TE LSP Count sub-TLV is OPTIONAL and MUST NOT appear more than once within the Link TLV (Type 2) that is itself carried within either the Traffic Engineering LSA specified in [RFC3630] or the OSPFv3 Intra-Area-TE LSA (function code 10) defined in [RFC5329]. If a second instance of the Unconstrained TE LSP Count sub-TLV is present, the receiving system MUST only process the first instance of the sub-TLV.

The OSPF Unconstrained TE LSP Count sub-TLV format is defined below:

Type (2 octets): 23

Length (2 octets): 4

Value (4 octets): number of unconstrained TE LSPs signalled across the link.

4. Elements of Procedure

The absence of the Unconstrained TE LSP Count sub-TLV SHOULD be interpreted as an absence of information about the link.

Similar to other MPLS Traffic Engineering link characteristics, LSA/LSP origination trigger mechanisms are outside the scope of this document. Care must be given to not trigger the systematic flooding of a new IS-IS LSP or OSPF LSA with a too high granularity in case of change in the number of unconstrained TE LSPs.

5. IANA Considerations

IANA has defined a sub-registry for the sub-TLVs carried in the IS-IS TLV 22 and has assigned a new TLV codepoint for the Unconstrained TE LSP Count sub-TLV carried within the TLV 22.

| Value | TLV Name | Reference |
|-------|--------------------------------------|-----------|
| 23 | Unconstrained TE LSP Count (sub-)TLV | RFC 5330 |

IANA has defined a sub-registry for the sub-TLVs carried in an OSPF TE Link TLV (type 2) and has assigned a new sub-TLV codepoint for the Unconstrained TE LSP Count sub-TLV carried within the TE Link TLV.

| Value | TLV Name | Reference |
|-------|--------------------------------------|-----------|
| 23 | Unconstrained TE LSP Count (sub-)TLV | RFC 5330 |

6. Security Considerations

The function described in this document does not create any new security issues for the OSPF and IS-IS protocols. Security considerations are covered in [RFC2328] and [RFC5340] for the base OSPF protocol and in [RFC1195] and [RFC5304] for IS-IS.

A security framework for MPLS and Generalized MPLS can be found in [G/MPLS].

7. Acknowledgements

The authors would like to thank Jean-Louis Le Roux, Adrian Farrel, Daniel King, Acee Lindem, Lou Berger, Attila Takacs, Pasi Eronen, Russ Housley, Tim Polk, and Loa Anderson for their useful inputs.

8. References

8.1. Normative References

- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [RFC5304] Li, T. and R. Atkinson, "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", RFC 5304, October 2008.
- [RFC5305] Li, T. and H. Smit, "IS-IS extensions for Traffic Engineering", RFC 5305, October 2008.
- [RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, September 2008.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

8.2. Informative References

- [G/MPLS] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", Work In Progress, July 2008.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

Authors' Addresses

JP Vasseur (editor)
Cisco Systems, Inc
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

EEmail: jpv@cisco.com

Matthew R. Meyer
BT
Boston, MA
USA

EEmail: matthew.meyer@bt.com

Kenji Kumaki
KDDI R&D Laboratories, Inc.
2-1-15 Ohara Fujimino
Saitama 356-8502, JAPAN

EEmail: ke-kumaki@kddi.com

Alberto Tempia Bonda
Telecom Italia
via G. Reiss Romoli 274
Torino, 10148
ITALIA

EEmail: alberto.tempiabonda@telecomitalia.it

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.